

# Mitigating Timing Side Channel in Shared Schedulers

Sachin Kadloor<sup>†\*</sup>, Negar Kiyavash<sup>‡\*</sup>, Parv Venkitasubramaniam<sup>§</sup>

<sup>†</sup> ECE Department and Coordinated Science Lab.

<sup>‡</sup> ISE Department and Coordinated Science Lab.

<sup>\*</sup>University of Illinois at Urbana-Champaign

<sup>§</sup>ECE department, Lehigh University

{kadloor1,kiyavash}@illinois.edu, parv.v@lehigh.edu

**Abstract**—In this work, we study information leakage in timing side channels that arise in the context of shared event schedulers. Consider two processes, one of them an innocuous process (referred to as Alice) and the other a malicious one (referred to as Bob), using a common scheduler to process their jobs. Based on when his jobs get processed, Bob wishes to learn about the pattern (size and timing) of jobs of Alice. Depending on the context, knowledge of this pattern could have serious implications on Alice’s privacy and security. For instance, shared routers can reveal traffic patterns, shared memory access can reveal cloud usage patterns, and suchlike. We present a formal framework to study the information leakage in shared resource schedulers using the pattern estimation error as a performance metric. The first-come-first-serve (FCFS) scheduling policy and time-division-multiple-access (TDMA) are identified as two extreme policies on the privacy metric, FCFS has the least, and TDMA has the highest. However, on performance based metrics, such as throughput and delay, it is well known that FCFS significantly outperforms TDMA. We then derive two parametrized policies, accumulate and serve, and proportional TDMA, which take two different approaches to offer a tunable trade-off between privacy and performance.

## I. INTRODUCTION

It has long been known that resources shared between processes lead to covert and side channels that can leak information from one process to another. A covert communication channel is one which is not normally intended to be used for communication [1], infact, its existence is usually unknown to the system designer. Covert channels are typically used by a trusted insider with access to a secret piece of information to convey it to an outsider. Examples of covert channels include, embedding information in the unused header files of network protocols [2], and two processes running on a computer communicating with each other through the access patterns of the shared memory [3]. In a covert channel, one process structures its use of the shared resource in a particular pattern so as to communicate secret information to another. Covert channels have been studied extensively in the context of multi-level secure systems, where they can be used to create forbidden information flows [4], [5].

In contrast to a covert channel, in a side channel, one process tries to learn something about the operation of another

without the latter’s cooperation. Side channels, therefore, focus on information that is leaked *incidentally* by a victim process, rather than explicitly coded by a sender. Examples of such channels include: an attacker non-invasively improving his odds of guessing cryptographic keys used by a crypto-system by observing its instantaneous power usage [6], making use of the fact that the power usage for a certain set of CPU operations is higher than others; an eavesdropper trying to guess the underlying communication by observing the encrypted packets flowing across a link [7], where he makes use of the fact that encryption does not alter the volume and timing of packets flowing on the link.

In this work, we consider the timing side channel that exists inside of a shared scheduler. A timing side channel is one in which information is conveyed (leaked) through the timings of various events. Schedulers are used in multi-tasking systems where they dictate how a finite resource is to be divided among several competing processes. Examples of such systems include: hardware resources (CPU, storage, buses) inside of a computer being shared among different processes, multiple network streams flowing through a common router, a cloud based shared computing infrastructure, etc.. In such systems, the *quality-of-service* experienced by one user of the system is directly influenced by the activities of the other users of the system. For example, a sudden slowdown in web access speeds of one user could indicate an increase in network usage from the other users sharing the same network infrastructure. In this manner, a shared scheduler incidentally creates a timing side channel through which a malicious user could potentially learn about the activities of the other users using the system.

For the remainder of this work, these systems are abstracted out as a processor being shared by multiple users, as shown in Figure 1. Jobs arrive from multiple users to the scheduler. The processor can work on only one job at a time. The scheduler queues up the incoming jobs, and dynamically decides on how the processor time gets divided amongst these competing jobs. In this manner, the delays experienced by jobs from one user are directly influenced by the number, timing, and size of the jobs issued by the other users.

Some of the commonly used schedulers are *first-come-first-served* (FCFS): jobs are served in the order they are issued to the scheduler, *time-division-multiple-access* (TDMA): each




Fig. 1. An abstraction of a shared scheduling system




Fig. 2. An event/packet scheduler being exploited by a malicious user to infer the arrival pattern from the other.

user is pre-assigned time slots during which the processor serves only jobs from that user, *round-robin* (RR): one job is served from each of the queued up users in succession, *priority schedulers*: one class of jobs are served ahead of another class of jobs according to some pre-defined rule, *shortest-job-first* (SJF): a particular type of priority scheduler where jobs that require smaller time to be processed are served first. Each of these schedulers is optimized to perform well on a different performance metric, such as, *throughput*: number of job completions per unit time, *average delay*: the mean time difference between the job completion and the job arrival, *fairness*: a metric to measure if the resource is being distributed equally/fairly between the processes, etc.. Like any engineered system, a system designer has to make a calculated trade-off among these conflicting metrics while picking a suitable scheduler. *We argue that while choosing a scheduler that serves jobs from multiple non-trusting users, one has to consider the privacy offered by it along with the other metrics.* In this work, we explore the resulting trade-offs one has to make if privacy is taken into account.

We consider the scenario when a scheduler is serving jobs from two users, where one of them is an innocuous user and other a malicious user. The malicious user, Bob, wishes to learn the pattern of jobs sent by the innocuous user, Alice. Bob exploits the fact that when the processor is busy serving jobs from Alice, his own jobs experience a delay. As shown in Figure 2, Bob computes the delays experienced by his jobs and uses these delays to infer about the times when Alice tried to access the processor, and possibly the sizes of jobs scheduled. Learning this traffic pattern from Alice can aid Bob in carrying out traffic analysis attacks.

A summary of our main contributions follow.

- 1) *Development of an analytical framework to characterize the privacy performance of a scheduling policy.* While evaluating the privacy performance of a scheduling policy, we consider the scenario described in Figure 2, Alice, the innocuous user, and Bob, the malicious user, are the only two users of the system. Arrivals from Alice are modeled as a Poisson process. Bob, on the other hand, is allowed to pick the time and size of the jobs

he issues to a scheduler. Furthermore, he is assumed to know the scheduling policy being used. The privacy offered by the policy is defined as the mean square error incurred by Bob when estimating Alice's arrival pattern when he picks an optimal attack strategy. Higher the error, the better the policy is at protecting the privacy of the users.

- 2) *Evaluation of the privacy metric of commonly deployed scheduling policies.* FCFS is one of the most commonly deployed scheduling policies owing to its simplicity. It is throughput optimal and results in minimal queuing delay. However, by the nature of the policy, there is a large correlation between the waiting times of jobs of one user and the arrival pattern of the other. Consequently, as shown in Section IV-C by the explicit construction of one attack strategy, FCFS is the weakest policy on the privacy metric. On the other hand, TDMA, wherein the delays experienced by jobs of one user are completely independent of the arrivals of the other, ranks highest on the privacy metric. However, TDMA is a highly inefficient policy in terms of throughput and delay, especially when the traffic is varying. It is especially inefficient when the number of users using the scheduler is large.
- 3) *Design policies that offer good privacy-delay trade-offs.* We design two parametric policies, *accumulate-and-serve* and *proportional-TDMA* which can be tuned to trade-off performance for improved privacy. These policies take two different approaches to achieve privacy, one pre-distorts timing information, other pre-allocates processor times to different users. Unlike TDMA, both these policies are throughput-optimal.

## II. RELATED WORKS

The current work was motivated by an earlier work of ours, [8], wherein, we demonstrated that a side channel does exist in DSL routers, wherein a network path is shared between all the incoming packet streams. The attack is briefly described below. The attacker sends equally spaced ping packets to Alice's DSL router from his home computer. He then observes the round trip times (RTTs) of the packets. The traffic entering Alice's computer, and Bob's RTTs are shown in Figure 3; as can be observed, there is a clear correlation between the two. The fact that DSL routers employed FCFS scheduling aided the attack. Such an attack gives the attacker a noisy observation of the timing and the sizes of packets entering Alice's computer. Although the contents of the packets are not revealed, learning such timing information opens up the possibility for the attacker to carry out remote traffic analysis. Some of the instances of traffic analysis include recovery of information about keystrokes typed [9], [10], websites visited [11], [12], or words spoken over VoIP [13]. In all these works, the attacker observes Alice's traffic and uses statistical inference techniques to carry out the attack. In [14], the authors consider the scenario where a client is connected to a rogue website using a TOR network, which is designed to protect

the identity of the users. The website modulates the traffic sent to the client. The website can then try to simultaneously send data through each of the TOR nodes and measure the delay incurred. By correlating this delay with the traffic it sent to the client, the website can obtain its identity, thus defeating the purpose of TOR. While that attack is no longer viable [15], the reason is that there are many more TOR nodes now than they were when [14] was published, and not because the timing based side channel has been eliminated. In [16], the authors exploit the side channel in a DSL router to infer the website being visited by the victim.

The increasing interest in cloud computing, wherein users issue jobs to a shared computing platform, opens up possibilities for such an attack, and is another motivation to study these timing side channels. In [17], the authors map the internal infrastructure of Amazon's EC2 cloud computing service, and demonstrate that it is possible for an attacker to place his virtual machine (VM) on the same physical computer as the target's VM. They shown that once placed on the same physical computer, any timing channel created by sharing of the processor can be exploited by the attacker. In cryptographic side-channels, the attacker aims at recovering cryptographic keys by utilizing the timing variations required for cryptographic operations [18], [19].

While timing covert channels have been studied extensively, most notably [20], there has been very little work in studying timing side channels. Most of the solutions proposed to mitigate the information leakage in side channels are system specific. Some examples are: the most common mitigation technique against such channels is cryptographic blinding [21], [19]; in the context of language-based security, Agat [22] introduces a program transformation to remove timing side channels from programs written in a sequential imperative programming language; the NRL Pump proposed for mitigating timing channels that arise in multilevel security systems (MLS) when a high confidentiality processes can communicate through Acks he sends to a low confidentiality processes [23].

The paper is organized as follows. In Section III, we formally introduce a system model, and the metric of performance that we use to compare the privacy of different scheduling policies. In Section IV-A, we quantify the highest degree of privacy that any scheduling policy can guarantee, and demonstrate that TDMA scheduling policy provides this. We start in Section IV-C discussing an attack strategy against FCFS policy to demonstrate that it does leak significant information between the flows. In Section V, we discuss two strategies for designing provably secure scheduling policies. Finally, in Section VI, we derive the mean delay experienced by a job when each of these policies is used, and comment on the delay-privacy tradeoff offered by the different policies.

### III. SYSTEM MODEL AND DEFINITIONS

In this section, we introduce formally a model of the scheduling system, and introduce a metric which measures the strength of the policy in preserving the privacy of the users. The scheduler is modeled as an infinite buffer server




Fig. 3. Real traffic on a DSL line vs. observed probe RTTs when Alice is browsing the website [www.yahoo.com](http://www.yahoo.com)

that is serving jobs from two users. The scheduler can serve jobs at a rate of one per unit time. We consider the scenario when one of them is an innocuous user and other a malicious one. However, the scheduler is unaware of who is malicious and who is innocuous; therefore any policy it picks should not distinguish between the users. The malicious user, Bob, wishes to exploit the queuing side channel described earlier to learn about the pattern of jobs sent by the innocuous user, Alice. Bob is assumed to know accurately the time when his jobs are issued, and the time it took for the scheduler to process it, i.e. the difference between the completion time of the job and the time when it was issued. Knowing the delays experienced by his jobs, Bob uses this information to guess the arrival pattern of jobs from Alice.

The ability of Bob to successfully learn about Alice's arrival process depends heavily on Alice's arrival process itself. For example, on-off patterns are easier to detect reliably compared to an arrival process that is less bursty. In order to ensure that the scheduling policies we design are robust to a variety of arrival patterns, Alice's arrival process will be modeled as a Poisson process of rate  $\lambda_2$ , with all the jobs of unit size. We do this partly because Poisson processes are known to have maximum entropy rate among processes of a given rate [24], and hence represent a rich class of arrival processes. Further, the analytical tractability of Poisson processes in a queuing system reveal the nature of fundamental trade-offs between privacy and delay in this system. We will comment on the case when Alice's traffic pattern follows a general arrival pattern later. Furthermore, we assume that Bob is aware of the statistical description of arrivals from Alice. A policy that guards the privacy of Alice in this scenario will also perform well when the attacker does not know a priori this rate.

### A. Measuring the strength of the scheduling policy: a privacy metric

Alice issues unit sized jobs to the scheduler according to a Poisson process of rate  $\lambda_2$ . The total number of jobs issued by Alice until time  $u$  is given by  $\mathcal{A}_A(u)$ . The malicious user, Bob, also referred to as the attacker, issues his jobs at times  $t_1^n \doteq \{t_1, t_2, \dots, t_n\}$ , and is free to choose their sizes,  $s_1^n \doteq \{s_1, s_2, \dots, s_n\}$ , as well.<sup>1</sup> Let  $t_1'^n \doteq \{t_1', t_2', \dots, t_n'\}$  be the departure times of these jobs. Bob makes use of the observations available to him, the set  $\{t_1^n, s_1^n, t_1'^n\}$  and the knowledge of the scheduling policy used, in estimating Alice's arrival pattern. The arrival pattern of Alice is the sequence  $\{X_k\}_{k=1,2,\dots,N}$ , where  $X_k = \mathcal{A}_A(kc) - \mathcal{A}_A((k-1)c)$ , is the number of jobs issued by Alice in the interval  $((k-1)c, kc]$ , referred to as the  $k^{\text{th}}$  clock period of duration  $c$ .  $Nc$  is the time horizon over which the attacker is interested in learning Alice's arrival pattern.

The privacy offered by a scheduling policy is measured by the long run estimation error incurred by Bob in such a scenario when he is free to choose the number of jobs he issues, times when he issues them and their sizes, subject to a maximum rate constraint, and when he optimally estimates Alice's arrival pattern. Formally, the privacy offered by a scheduling policy is defined to be:

$$\mathcal{E}_{\text{Scheduling policy}}^{c, \lambda_2} = \lim_{N \rightarrow \infty} \min_{\substack{n, t_1^n, s_1^n: \frac{\sum_{i=1}^n s_i}{Nc} < 1 - \lambda_2}} \sum_{k=1}^N \mathbf{E} \left[ \left( X_k - \mathbf{E} \left[ X_k | t_1^n, t_1'^n, s_1^n \right] \right)^2 \right], \quad (1)$$

where, the expectation is taken over the joint distribution of the arrival times of Alice's jobs, the arrival times and sizes of jobs from the attacker and his departure times. This joint distribution is in turn dependent on the scheduling policy used, which is known to the attacker. Finally, the attacker is assumed to know the statistical description of Alice's arrival process, and he is allowed to pick  $\sum_{i=1}^n s_i / Nc$ , the average rate at which he issues his jobs, to be any value that is less than  $1 - \lambda_2$ , so as to keep the system stable. A scheduling policy is said to preserve the privacy of its users if the resulting estimation error is high.

*A game-theoretic viewpoint of the privacy metric:* The optimal value of the estimation error can be viewed in a game theoretic framework as a saddle point between the scheduler and Bob; Bob's task is to minimize estimation error, and the scheduler's task is to maximize it. However, we force the scheduler to *play first* and announce the scheduling policy before the attacker chooses an attack strategy. This is motivated by practical considerations. In most system, the scheduling protocol is known to the users apriori. Further, in side channel attacks, it is not possible for the scheduler to

identify a malicious user and tailor the policy accordingly. Consequently, this assumption empowers the attacker and the achievable privacy would be a benchmark for attackers equipped with lesser information.

Our motivation for using the minimum mean squared error as a metric of performance is as follows. The minimum mean squared error, as considered in this paper, does not conform to a specific adversarial learning technique, but serves as a universal lower bound over all adversarial strategies taking into account the complete available information for the entire duration of the system operation. A natural alternative metric would be to measure the information leakage using Shannon's equivocation  $\lim_{N \rightarrow \infty} \frac{1}{N} \sum_{k=1}^N H(X_k | t_1^n, t_1'^n, s_1^n)$ . While entropy serves as a measure of uncertainty, which guarantees a minimum probability of error for an adversary (Fano's inequality), MMSE bounds the actual error incurred. The purpose of quantifying privacy is to have a meaningful measure of how breachable a system is, and in that respect, both these measures provide that interpretation. Furthermore, the two metrics are related in the sense that both MMSE and entropy are functionals of the probability mass function of the same conditional random variable,  $X_k | t_1^n, t_1'^n, s_1^n$ . MMSE is the variance of the random variable, while equivocation is the entropy of the random variable. It can be shown, using proof techniques similar to those in [26], that large estimation error does correspond to large entropy, although the relationship between the two is not monotonic.

In the subsequent section, we provide a bound on the maximum estimation error an attacker could incur, demonstrate that TDMA scheduling policy ensures that the attacker incurs this error. Subsequently, we present a scheduling policy that achieves an estimation error close to the maximum estimation error at the cost of a limited increase in delay.

## IV. PERFORMANCE OF FCFS AND TDMA ON THE PRIVACY METRIC

In this section, we prove that the FCFS and TDMA are two extreme policies on the privacy metric, FCFS offers the least privacy while TDMA offers the highest.

### A. An upper bound on the estimation error of the attacker

In order to gauge the amount of information that the attacker learns through this side channel, it is important to first study the amount of information that the attacker has even before performing any attack. In our case, the attacker is assumed to know the statistical description that governs the arrival pattern from Alice, that it is a Poisson traffic of rate  $\lambda_2$ . In this section, we compute the estimation error incurred by the attacker when he uses just this statistical description to estimate Alice's arrival pattern. This is the maximum estimation error that a rational attacker can incur, and will also serve as a benchmark to test the efficacy of a scheduling policy in preserving Alice's privacy.

*Theorem 4.1:* Irrespective of the scheduling policy used, the maximum estimation error that the attacker can incur is  $\lambda_2 c$ .

<sup>1</sup>We have also worked on a version of this problem where the attacker is also forced to issue jobs only of unit size, refer [25]. The present work is therefore a generalization of the former.

*Proof:* When Alice's traffic is a Poisson  $\lambda_2$  traffic, the number of arrivals in between two clock ticks,  $X_k$ , is a Poisson random variable with parameter  $\lambda_2 c$ . Also  $X_k$ s are independent and identically distributed (i.i.d.). Ignoring all the observations available to him, viz.  $\{t_1^n, t_1'^n, s_1^n\}$ , if Bob estimates  $X_k$  with its statistical mean,  $\lambda_2 c$ , for all  $k$ , the mean estimation error incurred by him is equal to the variance of the Poisson random variable, which is  $\lambda_2 c$ . ■

Define

$$\mathcal{E}_{\text{Max}}^{c, \lambda_2} \doteq \lambda_2 c. \quad (2)$$

$$\geq \mathcal{E}_{\text{Scheduling Policy}}^{c, \lambda_2} \quad (3)$$

Making a clever use of his observations, the attacker can potentially perform a better job at guessing Alice's arrival process, thereby incurring a lower estimation error. A good scheduling policy should have an estimation error as close  $\mathcal{E}_{\text{Max}}^{c, \lambda_2}$  as possible.

*B. Time Division Multiple Access (TDMA) policy achieves the maximum privacy*

If the scheduler uses a Time-Division-Multiple-Access (TDMA) policy, it allocates dedicated time slots to process jobs from each user. For instance, when there are only two users, the scheduler could assign the odd time slots to serve jobs from the first user, and even time slots to process jobs from the other. If there are no unserved jobs from one of the users, the scheduler just idles in the corresponding time slots.

For such a policy, the arrival pattern of one user does not influence the completion times of jobs in the other. In this scenario, the attacker can do nothing better than use his knowledge of the Alice's arrival statistics to guess the arrival pattern. Therefore, the estimation error incurred would be  $\mathcal{E}_{\text{TDMA}}^{c, \lambda_2} = \lambda_2 c$ . This means that from an privacy perspective, TDMA is an optimal albeit impractical scheduling policy.

The high privacy offered by TDMA comes at a significant performance cost in terms of throughput and delay. First, note that the policy is not throughput optimal [27], [28]. By throughput optimality, we mean the following. Suppose there are  $m$  users using the scheduler to process their jobs, and user  $i$  issues unit sized jobs at a rate  $\lambda_i$ , and the scheduler can process 1 unit of job in a unit interval of time. It can be shown that the system stays stable, i.e., the queue sizes do not blow to infinity, for any values of the input rates as long as they satisfy  $\sum_{i=1}^m \lambda_i < 1$  if the scheduler uses the FCFS policy.

However, if the scheduler used TDMA, then for the system to be stable, we would require  $\lambda_i < 1/m$ , which significantly shrinks the *rate region* over which the system stays stable. Second, even if the arrival rates from users are such that the system stays stable, as user's job patterns are rarely periodic, they would therefore incur severe delays. The delay issues are further discussed in Section VI.

*C. FCFS offers the least privacy*

The FCFS scheduling policy serves jobs in the order in which they arrive. Although this policy is easy to implement,

and fares well in the metrics of throughput and mean delay, as we will demonstrate, it does not perform well in terms of preserving the privacy of a user. For the system model considered here, the following theorem can be proved. It states that the estimation error incurred by the attacker when FCFS scheduling policy is used is equal to zero when the attacker issues his jobs at a high enough rate.

*Theorem 4.2:* FCFS policy offers no privacy to its users. Specifically, for a fixed arrival rate from Alice,  $\lambda_2$ , the estimation error incurred by the strongest attacker is equal to zero. That is,

$$\mathcal{E}_{\text{FCFS}}^{c, \lambda_2} = 0. \quad (4)$$

*Proof:* We prove this theorem by specifying one particular attack strategy, i.e., one set of arrival times and sizes of the jobs from the attacker which guarantees zero error. The attacker issues one job every  $c/\lceil c \rceil$  time units, where  $\lceil c \rceil$  is the smallest integer greater than or equal to  $c$ . The size of each job is  $\lambda c/\lceil c \rceil$ , so that the rate at which he issues jobs is equal to  $\lambda$ . Therefore  $t_k = kc/\lceil c \rceil$  and  $s_k = \lambda c/\lceil c \rceil$ . Let  $t_1^n$  be the departure times of his jobs, and  $\tilde{X}_k$  be the number of jobs issued by Alice in between times  $(k-1)c/\lceil c \rceil$  and  $kc/\lceil c \rceil$ . Recall that  $X_k$  is the number of jobs issued by Alice in between times  $(k-1)c$  and  $kc$ . Therefore,  $X_k = \sum_{i=1}^{\lceil c \rceil} \tilde{X}_{(k-1)\lceil c \rceil + i}$ . This

means if the attacker estimates the sequence  $\{\tilde{X}_k\}$  accurately, he can estimate the sequence  $\{X_k\}$  accurately as well. Note that the time between successive arrivals from the attacker is  $c/\lceil c \rceil < 1$ . The size of a job from Alice is 1. Hence the attacker can accurately learn whether Alice issued a job between two of his jobs or not. His estimation procedure is the follows.

Consider first the scenario when  $t_{k-1}' < t_k$ , that is, his  $k-1^{th}$  job departs before the arrival of  $k^{th}$  job. In this case, suppose  $\tilde{X}_k = 0$ , then  $t_k' = t_k + s_k$ , i.e., the  $k^{th}$  job goes into service immediately upon its arrival. This scenario is shown in Figure 4. On the other hand, if  $\tilde{X}_k > 0$ , shown in Figure 5 then  $t_k + s_k + \tilde{X}_k - 1 < t_k' < t_k + s_k + \tilde{X}_k$ , which implies  $\tilde{X}_k = \lceil t_k' - (t_k + s_k) \rceil$ .

Now, suppose  $t_{k-1}' > t_k$ , i.e., if the  $k-1^{th}$  job from the attacker departs after the arrival of his  $k^{th}$  job, shown in Figure 6. In this case,  $\tilde{X}_k = t_k' - s_k - t_{k-1}'$ .

Clearly,  $\tilde{X}_k$  is a deterministic function of  $t_{k-1}, t_{k-1}', t_k, t_k'$  and  $s_k$ . Therefore, the attacker incurs zero error in estimating Alice's arrival pattern. ■

Next we present some remarks about the result derived above. For the attack strategy specified in Theorem 4.2, the attacker incurs 0 error independent of his rate. This means that the attacker can estimate Alice's arrival pattern exactly even when he is permitted to issue jobs at an arbitrarily small rate. This is the case because Alice's jobs are of size 1 and as long as the attacker issues jobs at a frequency greater than 1, he can estimate Alice's arrivals exactly. Note that the size of jobs issued by him can be arbitrarily small.

As an aside, the result of the theorem can easily be extended to the case when the job sizes of Alice are of arbitrary sizes




Fig. 4. This is the scenario when  $t'_{k-1} < kc/\lceil c \rceil$  and  $\tilde{X}_k = 0$ . In this scenario,  $t'_k = t_k + s_k$ . In the figure, the solid upward and downward pointing arrows denote the arrival and departure times, respectively, of attacker's jobs.




Fig. 5. This is the scenario when  $t'_{k-1} < kc/\lceil c \rceil$  and  $\tilde{X}_k = 1 > 0$ . In this scenario,  $\tilde{X}_k = \lceil t'_k - (t_k + s_k) \rceil$ . In the figure, the solid upward and downward pointing arrows denote the arrival and departure times, respectively, of attacker's jobs, and the dotted arrows represent the jobs from Alice. The size of the arrow is proportional to the size of the job.

(for example, when the job sizes from Alice are exponentially distributed). The attack strategy in such a scenario would be the follows. The attacker issues one job every  $\delta$  units of time, and the size of each job is  $\lambda\delta$ . His estimate of Alice's arrivals is the follows. If  $t'_{k-1} < t_k$  and  $t'_k = t_k + s_k$ , then he estimates the number of arrivals from Alice between times  $t_{k-1}$  and  $t_k$  to be zero. If  $t'_{k-1} < t_k$  and  $t'_k > t_k + s_k$ , then he estimates the number of arrivals from Alice between times  $t_{k-1}$  and  $t_k$  to be  $t'_k - (t_k + s_k)$ . And finally if  $t'_{k-1} > t_k$ , then he estimates the number of arrivals to be  $t'_k - s_k - t'_{k-1}$ . Like before, his estimation is accurate in the third scenario. However, in the first two scenarios, he could incur some error. Nonetheless, the error is bounded by  $\delta^2$ . Therefore, by choosing  $\delta$  small enough, the attacker can incur close to zero estimation error.

This suggests two possible alterations to the FCFS policy which could lead to a higher estimation error for the attacker. First, the policy could limit the frequency at which a user can issue jobs to the system, by imposing a minimum inter-arrival time. This could be done by having a token-bucket filter [29]. Second, the scheduler could impose a restriction on the smallest size of a job, perhaps by idling for a short while after serving a small job. While these mechanisms look promising, it is easy to see that they hurt the performance of the system in terms of a smaller throughput region (if the scheduler starts idling) and/or adding to the overall delay. As we shall show in the following section, one can design secure scheduling policies without sacrificing any of the throughput region.

To sum up, FCFS offers very little privacy to its users. On the other hand, TDMA offers the highest levels of privacy, albeit at a high performance penalty (in terms of delay and




Fig. 6. This is the scenario when  $t'_{k-1} > kc/\lceil c \rceil$ . In this scenario,  $\tilde{X}_k = t'_k - t'_{k-1} - s_k$ . In the figure, the solid upward and downward pointing arrows denote the arrival and departure times, respectively, of attacker's jobs, and the dotted arrows represent the jobs from Alice. The size of the arrow is proportional to the size of the job.

throughput). In the following section, we propose a parametric scheduling policy which provides a controlled trade-off between delay and privacy.

## V. PROVABLY SECURE SCHEDULING POLICIES

In this section, we derive two policies that are resistant to any attacks that the attacker might possibly launch. The first policy, *accumulate and serve* pre-distorts the arriving traffic before serving them. This erases the fine grained timing information that the attacker can learn. The second policy takes a different approach. TDMA leaks no information of one user's arrivals to the other because slots are statically reserved ahead of time. The second policy, *proportional TDMA*, uses this same principle to guarantee minimal information leakage, however, it is designed so that these reservations alter adaptively at a slow rate so as to minimize delays incurred by the users. We discuss these two policies in the following.

### A. Accumulate and Serve Policy

The Accumulate and Serve policy is shown in Figure 7. In this policy, the scheduler accumulates all the incoming jobs in its buffer for a period of  $T$  time units. It then serves all the accumulated jobs from user 1, followed by all the accumulated jobs from user 2, followed by those from user 3, and so on. In the two user scenario discussed before, user 1 could be the attacker, or Alice. If serving all these jobs takes  $M$  units of time and if  $M < T$ , then the scheduler idles for the remaining  $T - M$  time before beginning to serve the accumulated jobs. The intuition for such a policy stems from the fact that buffering destroys the timing information of the arrival times of the jobs, thereby reducing the amount of information that can be extracted by the attacker. The attacker can therefore learn, at most, the total number of jobs that arrived from Alice in the accumulate period, and nothing more fine grained. Our subsequent analysis will make use of this fact to show that this policy provides guaranteed levels of privacy to Alice, irrespective of any attack that can be launched by Bob. Using Foster-Lyapunov stability theorem, it can be shown that the scheduling policy is stable when the sum of rates from all the users is less than 1 (refer Appendix B).




Fig. 7. Pictorial representation of the accumulate and serve policy. The switches close every  $T$  time units. All the accumulated jobs from user 1 are served followed by jobs from user 2. The two flows correspond to jobs issued by the two users.

## B. Analysis

*Theorem 5.1:* A lower bound on the privacy offered by the accumulate and serve policy is given by,

$$\begin{aligned} \mathcal{E}_{\text{AccServe}}^{c, \lambda_2} &\geq \lambda_2 c \left(1 - \frac{c}{T}\right)_+ \\ &= \mathcal{E}_{\text{Max}}^{c, \lambda_2} \left(1 - \frac{c}{T}\right)_+, \end{aligned} \quad (5)$$

*Proof:* A detailed proof is given in Appendix A. ■

Some remarks about the consequences of Theorem 5.1:

- The accumulate and serve policy guarantees a certain level of privacy regardless of the attacker strategy. Furthermore, if the accumulate period is chosen to be large enough so that  $\frac{c}{T} \ll 1$ , then the estimation error incurred by the attacker gets increasingly close to the maximum estimation error, which is what he would have incurred by just using the statistics of Alice's arrival. The price paid for choosing a very large value for  $T$  is of course the delay experienced by jobs. Delay analysis is carried out in Section VI. When  $T$  is chosen to be 5 times  $c$  or greater, the ratio is greater than 0.8, irrespective of the attacker's strategy.
- It is important to note that the error does not depend on the strategy employed by the attacker. In particular, it does not depend on the rate at which the attacker issues his jobs, which seems unusual. This is a consequence of the result in Lemma A.1, which says that when the attacker is given minimum additional information (by a genie) about total number of jobs in an accumulate period, his arrival and departure times carry no additional useful information for the purpose of estimation. The rate at which attacker issues his jobs, however, does determine the gap between true estimation error and the genie aided error.
- The bound  $\mathcal{E}_{\text{Max}}^c \left(1 - \frac{c}{T}\right)_+$  evaluates to 0 when  $c \geq T$ , which is not very useful. However, we are mostly interested in the scenario when  $T$  is chosen to be large enough to guarantee a degree of protection against a given value of  $c$ .

## C. Proportional TDMA (p-TDMA) Policy

As stated before, TDMA achieves the highest privacy because the times at which the attacker gets served is known to him a priori, he learns nothing by performing the attack. The p-TDMA policy builds upon this idea. In the two user

scenario, the policy divides time into slots, and at time 0, the policy statically assigns all odd numbered slots to user 1 and even numbered slots to user 2, just like TDMA. At the end of an *adaptation period*  $L$ , the policy computes the empirical rates at which the two users have been issuing jobs to the scheduler,  $\hat{\lambda}_1$  and  $\hat{\lambda}_2$ . Between times  $L$  and  $2L$ , each time slot is reserved for user 1 with probability  $\frac{\hat{\lambda}_1}{\hat{\lambda}_1 + \hat{\lambda}_2}$ , and user 2 with probability  $\frac{\hat{\lambda}_2}{\hat{\lambda}_1 + \hat{\lambda}_2}$ . At time  $2L$ , the scheduler recomputes the empirical rates, and uses the new rates to decide whether to reserve a slot to serve user 1 or user 2. The policy continues to update the empirical rates every  $L$  time units. The idea is that the policy eventually learns the true rates at which the users issue jobs, and allocates times in proportion to these rates, thus reducing the delays incurred.

*Theorem 5.2:*

$$\mathcal{E}_{\text{p-TDMA}}^{c, \lambda_2} \geq \mathcal{E}_{\text{Max}}^{c, \lambda_2} \left(1 - \frac{c}{L}\right)_+. \quad (6)$$

*Proof:* The analysis is very similar to that of the accumulate and serve policy. The policy depends only on the empirical rates at which jobs are issued to the scheduler. If the attacker is given the side information about how many jobs Alice has issued in each adaptation period, his arrival and departure times carry no further information and can be discarded. The result follows. ■

Some remarks about this policy follow:

- Unlike the accumulate and serve policy that trades off delay for privacy, the p-TDMA policy trades off adaptation time for privacy. Larger the value of the adaptation period  $L$ , longer it takes for the policy to learn the true arrival rates. However, for stationary arrival processes, the long run average delay is independent of the duration of the adaptation period. The reason for this is the follows. Irrespective of the value of  $L$ , after sufficient number of adaptation periods, the empirical rates converge to the true rates, and stay that way. Therefore, in the long run, the delay experienced by the jobs is independent of  $L$ . In theory, by choosing a large enough value of  $L$ , this policy is guaranteed to perform close to TDMA on the privacy metric without the excessive delay penalty.
- In a practical system, if the arrivals are not stationary, i.e., the statistical description of the arrival process changes with time, or if a user issues jobs only for a short duration which is insufficient for the policy to learn the true arrival rate from him, the delays incurred by the user can still be large. One would have to choose the value of  $L$  wisely to strike a balance between the performance for short flows, and the privacy it offers.

## VI. DELAY ANALYSIS

The delay experienced by the jobs would depend on the total number of users of the system and the specific patterns of arrivals from them. For purposes of delay comparison, we will consider a scenario where all the users of the system have time invariant arrival statistics. This would be the delay incurred when all users are legitimate innocuous users. This






Fig. 8. Privacy-delay trade-off offered by different policies.  $\frac{D_{FCFS}}{D_{Policy}}$  measures the performance of a policy in terms of delay offered by it, higher the ratio, the better.  $\frac{E_{c, \lambda_2}^{Policy}}{E_{c, \lambda_2}^{TDMA}}$  measures the privacy performance of the policy, higher the better. We consider the case when two users use the system, one user issues jobs at a rate 0.2 and the other at 0.45. By varying the parameters associated with each of the two derived policies, we obtain curves for the accumulate and serve and p-TDMA policies. The value of  $c$ , the clock period, is set to two.

is a reasonable assumption, as during the ‘normal’ mode of scheduler operation, when there are only legitimate users using the scheduler, the policy should offer as minimum delay to the jobs as possible. We will consider the case where there are  $M$  users of the system, traffic from user  $i$  being a Poisson traffic with parameter  $\lambda_i$ , size of all the jobs being equal to 1 time unit. Define  $\lambda = \sum_{i=1}^M \lambda_i$ . We will compute the average delay of a job (averaged over all the jobs from all the users) for each of the policies discussed here.

*Theorem 6.1:* The mean delay experienced by a job for FCFS, TDMA, accumulate and serve and p-TDMA scheduling policies are as follows:

$$D_{FCFS} = 1 + \frac{\lambda}{2(1-\lambda)}, \quad (7)$$

$$D_{TDMA} = 1 + \frac{M}{2} + \sum_{i=1}^M \frac{\lambda_i}{\lambda} \frac{\lambda_i M^2}{2(1-\lambda_i M)}, \quad (8)$$

$$D_{Acc\ Serve} \leq 1 + \frac{\lambda(T+1)}{2} + \frac{\lambda + T(1-\lambda)^2}{2(1-\lambda)} \mathbf{I}_{\{\lambda > \lambda_T^*\}} + \sqrt{\lambda T} \mathbf{I}_{\{\lambda < \lambda_T^*\}}, \quad (9)$$

$$D_{p-TDMA} = 1 + \frac{1}{2(1-\lambda)} + \frac{M-1}{1-\lambda}. \quad (10)$$

where,  $\lambda_T^* = \frac{2T+1-\sqrt{1+4T}}{2T}$ .

*Proof:* A detailed proof is given in Appendix C. ■

Some remarks about the mean delay computed:

- For FCFS, accumulate and serve, and p-TDMA, the mean delay is finite, equivalently the system is stable when the sum of the rates from all the users is less than 1, i.e.,

$\lambda = \sum_{i=1}^M \lambda_i < 1$ . However, for TDMA, the mean delay is finite only when the rate from each user is less than  $1/M$ , i.e., when  $\lambda_i < 1/M$ . Therefore, TDMA loses out on multiplexing gains.

- $\lim_{\lambda \rightarrow 1} \frac{D_{Acc\ Serve}}{D_{FCFS}} = 1$ , i.e., at high traffic loads, the mean delay offered by the accumulate and serve policy is equal to that when FCFS policy is used. In the other extreme,  $\lim_{\lambda \rightarrow 0} \frac{D_{Acc\ Serve}}{D_{FCFS}} = 1 + \frac{T}{2}$ , which can be quite large especially if  $T$  is chosen to be large. For TDMA, the same limit would be  $1 + \frac{M}{2}$ . For p-TDMA, it is  $1/2 + M$ . Depending on the values of  $M$  and  $T$ , either TDMA or accumulate and serve might offer the least delays.
- In Figure 8, we plot the privacy offered by each of these policies vs the inverse of the delay offered by them. We consider the case when two users use the scheduler, one of them issues jobs at a rate 0.45 and the other at rate 0.2. FCFS offers the least delay but also the least privacy, while TDMA offers the highest privacy, but at the expense of a significant performance loss. Accumulate and serve, as discussed earlier, is a parametrized policy that can trade-off privacy for delay. On the other hand, the p-TDMA policy can be tuned to achieve highest privacy without affecting the delay performance at all. However, longer the adaptation period, longer the arrival streams from the users have to be for the higher initial delays to average out.

## VII. CONCLUDING REMARKS

In this work, we discuss the timing side channel that arises naturally when two processes share a common resource. We consider a system where jobs from two users arrive at a processor. The task is to find a scheduling policy for the processor that minimizes information leakage about one user’s job arrival pattern to the other while not significantly increasing the average delay seen by the jobs. We use the estimation error to quantify the information leak, and use it to demonstrate that the first come first served scheduling policy leaks significant information. We also propose two scheduling policies that provide guaranteed levels of privacy. Do these policies continue to be secure if the arrivals were a general arrival process instead of being Poisson? To answer this question, note that the operation of both these derived policies depend on the number of arrivals in a batch (in one accumulate period, or one adaptation period). Therefore, as long as the arrival process from Alice is such that, given the total number of arrivals in an accumulate/adaptation period, there is a large uncertainty in the arrival times of each jobs, the derived policies will still fare well.

## REFERENCES

- [1] D. P. Company, *A Guide to Understanding Covert Channel Analysis of Trusted Systems*. DIANE Publishing Company, 1994.
- [2] S. J. Murdoch and S. Lewis, “Embedding covert channels into tcp/ip,” in *Proceedings of the 7th international conference on Information Hiding*, IH’05, (Berlin, Heidelberg), pp. 247–261, Springer-Verlag, 2005.



- [3] Z. Wang and R. B. Lee, "Covert and side channels due to processor architecture," in *Proceedings of the 22nd Annual Computer Security Applications Conference, ACSAC '06*, (Washington, DC, USA), pp. 473–482, IEEE Computer Society, 2006.
- [4] J. K. Millen, "Covert channel capacity," in *IEEE Symposium on Security and Privacy*, pp. 60–66, 1987.
- [5] J. C. Wray, "An analysis of covert timing channels," *Security and Privacy, IEEE Symposium on*, vol. 0, p. 2, 1991.
- [6] P. Kocher, J. Jaffe, and B. Jun, "Introduction to Differential Power Analysis and Related Attacks," 1998.
- [7] A. Hintz, "Fingerprinting Websites Using Traffic Analysis," *Privacy Enhancing Technologies*, pp. 171–178, 2002.
- [8] S. Kadloor, X. Gong, N. Kiyavash, T. Tezcan, and N. Borisov, "A low-cost side channel traffic analysis attack in packet networks," in *IEEE ICC 2010 - Communication and Information System Security Symposium*, (Cape Town, South Africa), 2010.
- [9] D. X. Song, D. Wagner, and X. Tian, "Timing analysis of keystrokes and SSH timing attacks," in *USENIX Security Symposium*, 2001.
- [10] K. Zhang and X. Wang, "Peeping Tom in the Neighborhood: Keystroke Eavesdropping on Multi-User Systems," in *USENIX Security*, 2009.
- [11] M. Liberatore and B. N. Levine, "Inferring the source of encrypted HTTP connections," in *CCS '06: Proceedings of the 13th ACM conference on Computer and communications security*, (New York, NY, USA), pp. 255–263, ACM Press, 2006.
- [12] G. Bissias, M. Liberatore, D. Jensen, and B. Levine, "Privacy vulnerabilities in encrypted HTTP streams," in *Privacy Enhancing Technologies*, pp. 1–11, 2006.
- [13] C. V. Wright, L. Ballard, S. E. Coull, F. Monrose, and G. M. Masson, "Spot me if you can: Uncovering spoken phrases in encrypted voip conversations," in *SP '08: Proceedings of the 2008 IEEE Symposium on Security and Privacy*, (Washington, DC, USA), pp. 35–49, IEEE Computer Society, 2008.
- [14] S. J. Murdoch and G. Danezis, "Low-cost traffic analysis of tor," in *Proceedings of the 2005 IEEE Symposium on Security and Privacy*, SP '05, (Washington, DC, USA), pp. 183–195, IEEE Computer Society, 2005.
- [15] N. S. Evans, R. Dingledine, and C. Grothoff, "A practical congestion attack on tor using long paths," in *Proceedings of the 18th conference on USENIX security symposium, SSYM'09*, (Berkeley, CA, USA), pp. 33–50, USENIX Association, 2009.
- [16] X. Gong, N. Borisov, N. Kiyavash, and N. Schear, "Website detection using remote traffic analysis," in *Privacy Enhancing Technologies*, 2012.
- [17] T. Ristenpart, E. Tromer, H. Shacham, and S. Savage, "Hey, you, get off of my cloud: exploring information leakage in third-party compute clouds," in *Proceedings of the 16th ACM conference on Computer and communications security, CCS '09*, (New York, NY, USA), pp. 199–212, ACM, 2009.
- [18] D. Brumley and D. Boneh, "Remote timing attacks are practical," *Computer Networks*, vol. 48, no. 5, pp. 701–716, 2005.
- [19] P. Kocher, "Timing attacks on implementations of Diffie-Hellman, RSA, DSS, and other systems," in *Advances in Cryptology CRYPTO96*, pp. 104–113, Springer, 1996.
- [20] J. Giles and B. Hajek, "An Information-Theoretic and Game-Theoretic Study of Timing Channels," *IEEE Transactions on Information Theory*, vol. 48, pp. 2455–2477, September 2002.
- [21] D. Chaum, "Blind signatures for untraceable payments," in *Advances in Cryptology: Proceedings of Crypto*, vol. 82, pp. 199–203, 1983.
- [22] J. Agat, "Transforming out timing leaks," in *Proceedings of the 27th ACM SIGPLAN-SIGACT symposium on Principles of programming languages*, pp. 40–53, ACM, 2000.
- [23] I. S. Moskowitz and A. R. Miller, "The channel capacity of a certain noisy timing channel," *IEEE Trans. Inform. Theory*, vol. 38, pp. 1339–1344, July 1992.
- [24] J. A. McFadden, "The entropy of a point process," *SIAM Journal on Applied Mathematics*, vol. 13, no. 4, pp. 988–994, 1965.
- [25] S. Kadloor, N. Kiyavash, and P. Venkatasubramanian, "Mitigating timing based information leakage in shared schedulers," in *INFOCOM, 2012 Proceedings IEEE*, pp. 1044–1052, march 2012.
- [26] M. Feder and N. Merhav, "Relations between entropy and error probability," *Information Theory, IEEE Transactions on*, vol. 40, pp. 259–266, jan 1994.
- [27] R. Rom and M. Sidi, *Multiple Access Protocols: Performance and Analysis*. New York: Springer Verlag, 1990.
- [28] D. P. Bertsekas and R. G. Gallager, *Data Networks*. Englewood Cliffs, N.J.: Prentice-Hall, Inc., 1987.
- [29] D. D. Clark, S. Shenker, and L. Zhang, "Supporting real-time applications in an integrated services packet network: architecture and mechanism," *SIGCOMM Comput. Commun. Rev.*, vol. 22, pp. 14–26, Oct. 1992.
- [30] R. G. Gallager, *Poisson Processes, class notes*, available online at [http://www.rle.mit.edu/rgallager/documents/6.262ch2Dec10\\_000.pdf](http://www.rle.mit.edu/rgallager/documents/6.262ch2Dec10_000.pdf).
- [31] S. Asmussen, *Applied probability and queues*, vol. 5. John Wiley & Sons, Ltd., 1989.
- [32] S. Lam, "Delay analysis of a time division multiple access (tdma) channel," *Communications, IEEE Transactions on*, vol. 25, pp. 1489–1494, dec 1977.

## APPENDIX A PROOF OF THEOREM 5.1

Suppose that  $n$  jobs from Bob arrive in  $K$  accumulate periods. Let  $\{B_1, B_2, \dots, B_K\}$ , denoted in short by  $B_1^K$ , be the number of jobs arriving from Alice in accumulate periods  $1, 2, \dots, K$  respectively.  $B_1^K$  is the side information given to the attacker. The following lemma shows that the side information is the maximum information the attacker can hope to learn.

**Lemma A.1:** If the attacker is given the side information  $B_1^K$ , the arrival and completion times of his jobs carry no information and can be discarded. Let  $X_k$  be the random variable denoting the number of jobs that arrive from Alice in between clock ticks  $k-1$  and  $k$ . The conditional expectation,  $\mathbf{E}(X_k | t_1^n, t_1'^n, s_1^n, B_1^K)$  is the same as  $\mathbf{E}(X_k | B_1^K)$ .

**Proof:** Consider the conditional probability distribution  $\Pr(X_k | B_1^K, t_1^n, t_1'^n, s_1^n)$ . If accumulate and serve policy is used at the scheduler, the departure times of Bob's jobs are just a function of the arrival times and the size of his jobs, and the sequence  $\{B_1, B_2, \dots, B_K\}$ . Therefore,

$$\begin{aligned} \Pr(X_k | B_1^K, t_1^n, t_1'^n, s_1^n) &= \frac{\Pr(X_k, B_1^K, t_1^n, t_1'^n, s_1^n)}{\Pr(B_1^K, t_1^n, t_1'^n, s_1^n)} \\ &= \frac{\Pr(X_k, B_1^K, t_1^n, s_1^n)}{\Pr(B_1^K, t_1^n, s_1^n)} \frac{\Pr(t_1'^n | X_k, B_1^K, t_1^n, s_1^n)}{\Pr(t_1'^n | B_1^K, t_1^n, s_1^n)} \\ &= \Pr(X_k | B_1^K, t_1^n, s_1^n), \end{aligned} \quad (11)$$

where (11) follows because  $t_1'^n$  is a function of  $t_1^n, s_1^n$  and  $B_1^K$ . As a consequence, the three sets of random variables  $\{t_1^n, s_1^n\} - B_1^K - \{X_k\}$  form a Markov chain. Thus,  $\Pr(X_k | B_1^K, t_1^n, s_1^n) = \Pr(X_k | B_1^K)$ . ■

As a result of Lemma A.1, if the attacker is given  $B_1^K$ , the conditional distribution of  $X_k$ , and hence the resulting estimation error he will incur, is independent of the timing and number of jobs that the attacker sends. Also, if the resulting estimation error is denoted by  $\mathcal{E}_{\text{AccServe}}^{c, B_1^K}$ , then, from (??),

$$\mathcal{E}_{\text{AccServe}}^{c, \lambda} \geq \mathcal{E}_{\text{AccServe}}^{c, B_1^K}. \quad (12)$$

This is because, we are giving the attacker extra information which is not directly available to him. With this information, the attacker can only be better off in his estimation procedure. Hence,  $\mathcal{E}_{\text{AccServe}}^{c, B_1^K}$  bounds the estimation error that the attacker

incurs, and we will shortly show that this is large enough. Next, we turn our attention to computing  $\mathcal{E}_{\text{AccServe}}^{c, B_1^K}$ .

We first consider the case when the following hold:

- C1** The accumulate period,  $T$ , is an integer multiple of  $c$ , the duration of a clock period.
- C2** At time 0, the first clock tick aligns with the start of the first accumulate period.

If conditions **C1** and **C2** are met, every clock period is then completely contained within an accumulate period.

Suppose that the clock period  $k$  is contained within the accumulate period  $m$ . Let  $X_k$  denote the number of arrivals from Alice in clock period  $k$ , and let  $\bar{X}_k$  denote the number of arrivals from Alice in the other clock periods in accumulate period  $m$ . Then, it can be shown that, conditioned on the total number of arrivals in accumulate period  $m$ ,  $B_m$ ,  $X_k$  is independent of the arrivals in the other accumulate periods. Therefore,  $\Pr(X_k | B_1^K) = \Pr(X_k | B_m)$ . For a Poisson process, conditioned on the total number of arrivals in a given period, the arrival times themselves are uniformly distributed in the period [30]. Therefore, conditioned on  $B_m = \beta_m$ , the distribution of  $X_k$  is a Binomial random variable with parameters  $(\beta_m, \frac{c}{T})$ . Given the value of  $\beta_m$ , the best estimate for  $X_k$  is then  $\frac{c}{T}\beta_m$ . Thus, the estimation error incurred by the attacker in a slot in which  $\beta_m$  jobs arrived is equal to the variance of the Binomial random variable equal to  $\mathcal{E}_{\beta_m} = \beta_m \frac{c}{T} (1 - \frac{c}{T})$ . Given that  $\beta_m$  itself is a random variable, we need to average over its distribution to get the estimation error incurred by the attacker.  $\beta_m$  is Poisson with mean  $\lambda_2 T$ , and the estimation error in a typical clock period is

$$\begin{aligned} \mathbf{E}_{\beta_m}[\mathcal{E}_{\beta_m}] &= \mathbf{E}_{\beta_m}[\beta_m \frac{c}{T} (1 - \frac{c}{T})] \\ &= \lambda_2 c (1 - \frac{c}{T}) \\ &= \mathcal{E}_{Max}^c (1 - \frac{c}{T}), \end{aligned} \quad (13)$$

where (13) follows from (2).

When the conditions **C1** and **C2** do not hold, it can be shown the estimation error incurred by the attacker is greater than  $\mathcal{E}_{Max}^c (1 - \frac{c}{T})$ . It can be shown that the empirical average of the estimation error incurred in a clock period, as given in (13) is also equal to the time average of the estimation error incurred by the attacker, which, as defined in (1) is the privacy metric we are interested in. Therefore,  $\mathcal{E}_{\text{AccServe}}^{c, B_1^K} = \mathcal{E}_{Max}^c (1 - \frac{c}{T})$ .

*The case when  $c > T$ :* If  $c > T$ , then the side information  $B_1^K$  of the attacker is finer than the information which he wishes to extract from Alice's arrival process. Therefore, in such a case, the estimation error incurred by the attacker is equal to zero. Recall that in reality the attacker does not have this side information, and he has to infer it based on the arrival and departure times of his jobs. We are not interested in analyzing this case as we expect that in a system design, the accumulate period will be chosen to be sufficiently larger than  $c$ .

## APPENDIX B

### STABILITY OF THE ACCUMULATE AND SERVE POLICY

Let  $Q_n$  be the number of jobs waiting to be served at the beginning of the  $n^{\text{th}}$  accumulate period, i.e., at time  $nT^-$ . Because  $T$  is an integer multiple of the service time, note that there are no jobs with partially completed service at this time, and therefore the total work in the system at the end of accumulate period is always an integer. Let  $A_n$  denote the total number of jobs that arrive from all the users in the  $n^{\text{th}}$  accumulate period. Then,  $Q_n$  is a Markov chain with the state space  $\mathcal{Z}^+$ , and state update equation given by

$$Q_{n+1} = (Q_n + A_n - T)_+, \quad (14)$$

where the notation  $(i)_+$  stands for  $\max\{0, i\}$ . This is because, along with the jobs already waiting to be served at time  $nT^-$ ,  $A_n$  number of jobs arrive. Out of these, at most  $T$  of them get served in the accumulate period. The stability of this queue also implies that the mean waiting times of the jobs is always bounded as long as the sum of arrival rates from all the users is less than the service rate.

For a state  $q \in \mathcal{Z}^+$ , define the Lyapunov function to be  $V(q) = q^2/2$ . We will use Foster-Lyapunov's theorem to show that the Markov chain is positive recurrent. Let  $\lambda$  be the sum of the arrival rates. Consider the case when each of the arrival process is a Poisson process, then the cumulative arrival process is a Poisson process as well.  $A_n$  is then a Poisson random variable with parameter  $\lambda T$ . Given  $Q_n = q$ , the expected value of the drift is given by

$$\begin{aligned} \mathbf{E}[V(Q_{n+1}) - V(Q_n) | Q_n = q] &= \mathbf{E}\left[\frac{((q + A_n - T)_+)^2 - q^2}{2}\right] \\ &\leq \frac{1}{2} \mathbf{E}[(q + A_n - \alpha)^2 - q^2], \quad \alpha \leq T \quad (15) \\ &= \frac{1}{2} \left( (q - \alpha)^2 + \lambda T + (\lambda T)^2 + 2\lambda T(q - \alpha) - q^2 \right) \\ &= \underbrace{\frac{\lambda T + (\alpha - \lambda T)^2}{2}}_{K_1} - \underbrace{(\alpha - \lambda T)q}_{K_2}, \quad (16) \end{aligned}$$

where (15) follows because, for real numbers  $i, \alpha, \beta$ ,  $((i - \beta)_+)^2 \leq (i - \alpha)^2, \forall \alpha \leq \beta$ . By setting  $\alpha$  to  $T$ , the condition for the system to be stable translates to  $\lambda < 1$ . This is because, if  $\lambda < 1$ , then  $K_2 > 0$ , and therefore the expected drift of the Lyapunov function is negative for  $q$  large enough.

Furthermore, one can bound the mean number of jobs in the queue in the steady state by

$$\mathbf{E}[Q_n] \leq \frac{K_1}{K_2} = \frac{\lambda T + (\alpha - \lambda T)^2}{2(\alpha - \lambda T)}. \quad (17)$$

The bound in (17) holds for all values of  $\alpha$  in the range  $(\lambda T, T]$ . Choosing the value of  $\alpha$  that results in the tightest bound, it can be shown that

$$\mathbf{E}[Q_n] \leq \begin{cases} \sqrt{\lambda T} & \text{if } \lambda < \lambda_T^* \\ \frac{\lambda + (1 - \lambda)^2}{2(1 - \lambda)} & \text{if } \lambda \geq \lambda_T^* \end{cases}, \quad (18)$$

where  $\lambda_T^* = \frac{2T+1-\sqrt{1+4T}}{2T}$ .

## APPENDIX C PROOF OF THEOREM 6.1

The FCFS system is a simple  $M/D/1$  queue, and the mean delay can be derived from Polleczech and Khinchine's formula for the  $M/G/1$  system, refer [31]. For the proof for TDMA, refer [32]. The proof for accumulate and serve and the p-TDMA policies follow.

### A. Mean delay of accumulate and serve policy

The arriving jobs are stored in the buffer till the start of the next accumulate period. A job, chosen at random gets delayed by  $T/2$  time units at this stage. From the point of view of the processor, the inputs to the queue are batch arrivals which arrive every  $T$  time units, and the size of a batch is a Poisson random variable. The average delay across all jobs does not depend on the order in which jobs are served by the server, as long as the server does not idle. For delay analysis, we can therefore assume that the server serves these jobs in FCFS manner. In reality, recall that the server serves all the jobs from one user followed by all the jobs from the other user.

Let  $A_k$  denote the size of the batch that arrives at the end of the accumulate period  $k$  and  $Q_k$  be the number of jobs that are not yet served at the end of period  $k$ . By the FCFS assumption,  $Q_k$  is also the time the arriving batch waits before it gets served. The queue update equation is given in (14), and consequently, the bounds derived in Appendix B hold. Therefore, the mean waiting time before a batch of jobs starts getting served is upper bounded by  $\frac{\lambda+(1-\lambda)^2}{2(1-\lambda)}\mathbf{I}_{\{\lambda \geq \lambda_T^*\}} + \sqrt{\lambda T}\mathbf{I}_{\{\lambda < \lambda_T^*\}}$ .

In order to compute the mean delay experienced by a job, first note that a job drawn at random from the first  $K$  jobs, where  $K$  is a large number, has a higher chance of belonging to a bigger batch than a smaller one. In fact, the probability that a job drawn at random belongs to a batch of size  $b$  is given by  $\frac{b\mathbf{P}(A_k=b)}{\mathbf{E}[A_k]}$ . Given that the job is from a batch of size  $b$ , the mean number of jobs ahead of it can be shown to be  $\frac{b-1}{2}$ , consequently, conditional on the batch size to be  $b$ , the job has to wait for an additional time of  $\frac{b-1}{2}$  time units after the batch starts service before it can get served. Averaging over  $b$ , we can get the additional waiting time of a jobs drawn at random before it gets served to be

$$\sum_b b \frac{\mathbf{P}(A_k=b)}{\mathbf{E}[A_k]} \left( \frac{b-1}{2} \right) = \frac{\mathbf{E}[A_k^2] - \mathbf{E}[A_k]}{2\mathbf{E}[A_k]} = \frac{\lambda T}{2} \quad (19)$$

Once it gets to service, the service time of the job is a fixed 1 time unit. Therefore, the delay experienced by a job drawn at random, which is the mean delay of offered by the scheduling policy is bounded by (9).

### B. Mean delay of p-TDMA policy

If the arrival process from each user is a Poisson process of rate  $\lambda_i$ , in the steady state, the empirical arrival rates converge to the true arrival rates. User  $i$  gets service in a time slot with probability  $\lambda_i/\lambda$ . Let  $Q_n$  be the number of unserved

jobs from user  $i$  at the beginning of time slot  $n$ . We then have the following queuing equation,

$$Q_{n+1} = (Q_n - D_n)_+ + A_n,$$

where  $A_n$  denotes the number of job arrivals between times  $n$  and  $n+1$ , which is a Poisson random variable with mean  $\lambda_i$ , and  $D_n$  is a Bernoulli random variable which takes the value 1 if user  $i$  gets served in time slot  $n$ , which happens with probability  $\lambda_i/\lambda$ , or 0 otherwise. The queuing equation holds because, the system serves only jobs that have already arrived by time  $n$  in time slot  $n$ .

In the steady state, both  $Q_n$  and  $Q_{n+1}$  have the same distribution. Let  $q_i \doteq \Pr(Q_n = i) = \Pr(Q_{n+1} = i)$ , and  $\mathcal{Q}(z) \doteq \sum_{i=0}^{\infty} q_i z^i$  be the Z-transform of the steady state distribution. We then have

$$\mathcal{Q}(z) = \mathbf{E}[z^{Q_{n+1}}] = \mathbf{E}[z^{(Q_n - D_n)_+ + A_n}] \quad (20)$$

$$= q_0 \mathbf{E}[z^{A_n}] \left( q_0 \mathbf{E}[z^{(-D_n)_+}] + \sum_{i=1}^{\infty} q_i \mathbf{E}[z^{i-D_n}] \right) \\ = \frac{q_0 (\lambda_i/\lambda) e^{\lambda_i(z-1)} (1 - z^{-1})}{1 - (1 - (\lambda_i/\lambda) + (\lambda_i/\lambda) z^{-1}) e^{\lambda_i(z-1)}}, \quad (21)$$

where (20) follows from the queuing equation. Use the fact that  $\lim_{z \rightarrow 1} \mathcal{Q}(z) = 1$  to solve for  $q_0 = (1 - \lambda)$ . The mean number of unserved jobs from user  $i$  in the steady state is then equal to  $\sum_{i=0}^{\infty} i q_i = \lim_{z \rightarrow 1} z \mathcal{Q}'(z)$ , which can be computed. Using Little's law to relate the average queue length to the mean delay experienced by a job from user  $i$ , and averaging it across all users of the system, we get the result given in equation (10). The result also includes the time  $1/2$ , the mean time an arriving job waits before the nearest multiple of one time unit.